

A new combinatorial class of 3–manifold triangulations

Feng Luo and Stephan Tillmann

Abstract We define a new combinatorial class of triangulations of closed 3–manifolds, satisfying a weak version of 0–efficiency combined with a weak version of minimality, and study them using twisted squares. As an application, we obtain strong restrictions on the topology of a 3–manifold from the existence of non-smooth maxima of the volume function on the space of circle-valued angle structures.

AMS Classification 57M25, 57N10

Keywords 3–manifold, triangulation, circle-valued angle structure

1 Introduction

In computational topology, it is a difficult problem to certify that a given triangulation of a closed, irreducible, orientable 3–manifold is minimal. Here, the term *triangulation* includes semi-simplicial or singular triangulations. The difficulty mainly stems from the fact that on the one hand, the current computer generated censuses are limited in size (see [10, 3]), and on the other hand, it is difficult to find good lower bounds for the minimal number of tetrahedra (see [9, 6]). For many algorithms using 3–manifold triangulations, the 0–efficient triangulations due to Jaco and Rubinstein [4] have been established as a suitable platform. However, certifying that a given triangulation is 0–efficient involves techniques from linear programming.

This paper introduces a new class of triangulations satisfying a weak version of 0–efficiency combined with a weak version of minimality. The weak version of minimality (*face-pair-reduced*) is that certain simplification moves, which can be determined from the 2–skeleton, are not possible. The weak version of 0–efficiency (*face-generic*) puts mild restrictions on the combinatorics of the 2–skeleton. We show that a 0–efficient or minimal triangulation of a closed, irreducible, orientable 3–manifold is face-generic unless the underlying manifold is one of S^3 , $\mathbb{R}P^3$, $L(3, 1)$, $L(4, 1)$, $L(5, 1)$ or $L(5, 2)$. Moreover, if a triangulation is not face-generic and face-pair-reduced, then it can either be simplified to a face-generic, face-pair-reduced triangulation or it is easy to recognise which of the aforementioned six exceptions is the underlying manifold of the triangulation (Proposition 13).

The new notions, face-pair-reduced and face-generic, are motivated by key combinatorial properties of the 2–skeleton of minimal and 0–efficient triangulations respectively. They are the subject of Section 2. A new tool in the study of triangulations is introduced in Section 3: the study of (possibly pinched or immersed) surfaces that have a cell structure consisting of up to two squares, each of which is mapped to a *twisted square* in a tetrahedron. Twisted squares are also used to give a classification of the possible combinatorial types of the (possibly singular) 3–simplices in a face-pair-reduced, face-generic triangulation. In this paper, as is common usage, we call a singular 3–simplex in M a tetrahedron.

We give a key application to highlight the utility of our techniques. Previous work of the first author [8] defined and studied circle-valued angle structures on orientable, closed 3-manifolds using a volume function. It was shown that a smooth maximum of the volume function corresponds to a solution of the *generalised Thurston equation*, whilst a non-smooth maximum corresponds to a *cluster of three 2-quad type solutions* to Haken's normal surface equations. Such a cluster consists of three algebraic solutions with at most two non-zero quadrilateral coordinates, such that the set of their non-zero coordinates includes the three quadrilaterals supported by some tetrahedron. The present paper shows that (under sufficiently strong hypotheses on the triangulation), the existence of a non-smooth maximum gives strong information about the topology of the manifold. The following is a consequence of the more technical Theorem 23, which is stated and proved in Section 5:

Theorem 1 *Suppose M is a closed, orientable 3-manifold with a triangulation which is either minimal or (face-pair-reduced and face-generic). If there is a cluster of three 2-quad type solutions to Haken's normal surface equations, then either*

- (1) M is reducible, or
- (2) M is toroidal, or
- (3) M is Seifert fibred, or
- (4) M contains a non-orientable surface of genus 3.

A discussion of this theorem and its consequences can be found in [7] (§1 and §4.3), where it was first announced for minimal triangulations. In particular, Thurston's orbifold theorem implies that if M does not fall in the first three cases and contains a non-orientable surface of genus 3, then M satisfies Thurston's Geometrisation Conjecture without appeal to Perelman's Ricci flow techniques. The proof of the theorem shows that, for a 3-manifold with face-pair-reduced and face-generic triangulation, it is possible to deduce from the combinatorics of the triangulation which of the (not mutually exclusive) cases occurs (see Lemma 19, Proposition 20 and Theorem 23). As pointed out by the referee, it is worth mentioning that the main application of this paper can be viewed as a step towards realising a variant of Casson's program to prove the Geometrisation Theorem for 3-manifolds. Casson's program originally aimed to reprove Thurston's theorem via ideal triangulations in the case of 3-manifolds with boundary consisting of tori, and this program was later extended to closed 3-manifolds by the first author.

We have developed the new tools and notions for triangulations with a view towards keeping this paper brief, and the details at bay. They can be modified for the study of reducible 3-manifolds. Moreover, there is much scope in developing a full theory of face-pair-reduced and face-generic triangulations of irreducible 3-manifolds as they share many combinatorial properties with 0-efficient and minimal triangulations. This would, in particular, give new information about minimal triangulations of closed hyperbolic 3-manifolds. These remain elusive at the time of writing: they are sparse in the current censuses and no infinite families are known.

Acknowledgements The first author is partially supported by the USA National Science Foundation (project numbers DMS 1105808 and DMS 1222663). The second author is partially supported under the Australian Research Council's Discovery funding scheme (project number DP110101104), and thanks the Max Planck Institute for Mathematics, where parts of this work have been carried out, for its hospitality. The authors would like to thank the referee for constructive feedback.

2 Faces in triangulations

The properties of triangulations that we will need are a weak version of 0–efficiency combined with a weak version of minimality. The weak version of minimality (*face-pair-reduced*) is that certain simplification moves, which can be determined from the 2–skeleton, are not possible. The weak version of 0–efficiency (*face-generic*) puts mild restrictions on the combinatorics of the 2–skeleton.

2.1 Triangulations and normal surfaces

We will assume that the reader is familiar with standard definitions, facts and notions about triangulations of 3–manifolds and normal surfaces. Our main references are [4] for triangulations of 3–manifolds and [12] for normal surface theory. A triangulation of a compact, orientable 3–manifold M consists of a union of pairwise disjoint 3–simplices, $\tilde{\Delta}$, a set of face pairings, Φ , and a natural quotient map $p: \tilde{\Delta} \rightarrow \tilde{\Delta}/\Phi = M$. Since the quotient map is injective on the interior of each 3–simplex, we refer to the image of a 3–simplex in M as a *tetrahedron* and to its faces, edges and vertices with respect to the pre-image. Similarly for images of 2– and 1–simplices, which will be referred to as *faces* and *edges* in M .

A triangulation of the closed, orientable, connected 3–manifold M is *minimal* if M has no other triangulation with fewer tetrahedra. A triangulation of M is *0–efficient* if the only embedded, normal 2–spheres are vertex linking (see [4] or Section 4 for a definition of normal surface).

2.2 Small triangulations

If the closed, irreducible, orientable 3–manifold M has a (not necessarily minimal) triangulation with just one or two tetrahedra, then M is one of the lens spaces S^3 , $\mathbb{R}P^3$, $L(3, 1)$, $L(4, 1)$, $L(5, 1)$, $L(5, 2)$, $L(7, 2)$, $L(8, 3)$ or quaternionic space S^3/Q_8 (see, for instance, the census in [3]). Given such a small triangulation, one can recognise M from the combinatorics of the triangulation or (up to the ambiguity $L(5, 1)$ versus $L(5, 2)$) from $H_1(M)$.

A subset of the above list of manifolds of small complexity turns out to play a special role in the analysis of the combinatorics of the 2–skeleton of a triangulation. For simplicity, we will write $L(p)$ for a lens space with fundamental group \mathbb{Z}_p ; this mainly saves some space that would be taken up with distinguishing $L(5, 1)$ from $L(5, 2)$.

2.3 Faces in triangulations

There are eight types of faces in a triangulation, depending on how vertices or edges are identified (see Figure 1 in [4]). Ignoring possible identification of vertices, this reduces to the following five types based on the edge identifications only (see Figure 1 below). A face with no edge-identifications is termed a *triangle face*; it is a 2–simplex with some or all vertices identified. If a pair of edges is identified, the face is either a cone (possibly with the tip identified with a point on the boundary) and called a *cone face* or it is a Möbius band and called a *Möbius face*. In a Möbius face, we distinguish the *boundary edge* and the *core edge*. If all three edges are identified, the face is either a 3–fold or a dunce hat, and called a *3-fold face* or *dunce face* respectively.

The following results concerning minimal, 0–efficient or arbitrary triangulations are well-known:

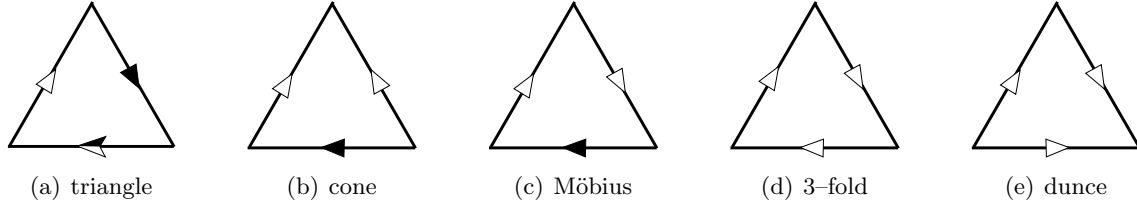


Figure 1: Types of faces. The Möbius, 3-fold and dunce faces have only got one vertex due to the edge identifications. A cone face may have one or two vertices and a triangle face may have up to 3 vertices.

Theorem 2 ([4], Theorem 6.1) *A minimal triangulation of the closed, irreducible, orientable 3-manifold M is 0-efficient unless $M = \mathbb{R}P^3$ or $L(3, 1)$.*

Proposition 3 *Suppose M is a closed, orientable 3-manifold with a 0-efficient triangulation.*

- (1) M irreducible and not $\mathbb{R}P^3$. ([4] Proposition 5.1)
- (2) The triangulation has one vertex or $M = S^3$ and there are precisely 2 vertices. ([4] Proposition 5.1)
- (3) If an edge bounds an embedded disc, then $M = S^3$. ([4], Proposition 5.3)
- (4) No face is a cone unless $M = S^3$. ([4], Corollary 5.4)

Proposition 4 *Suppose M is a closed, orientable, irreducible 3-manifold with an arbitrary triangulation. If a face in the triangulation is a 3-fold, then $M = L(3)$.*

Proof Let N be a regular neighbourhood of the 3-fold F in M . Then $2 = 2\chi(F) = 2\chi(N) = \chi(\partial N)$. If ∂N is connected, then it is a sphere and the result follows since M is irreducible and the 3-fold is then a spine for M . To see that ∂N must be connected, equip the interior of the triangle that makes up the 3-fold with a transverse orientation. The single edge in the boundary of the 3-fold is a simple closed curve γ in M . Let $p \in \gamma$ and B_p be a small regular neighbourhood of p . Then $F \cap B_p$ looks like the shape of a Y times an interval, and there must be at least one component of $B_p \setminus (F \cap B_p)$ with the property that the transverse orientation points both into the component and out of it. But this implies that F cannot separate N , whence N has connected boundary. ■

The starting point for our study of triangulations is the following:

Lemma 5 *Suppose M is a closed, orientable, irreducible 3-manifold with an arbitrary triangulation, and σ a tetrahedron in M . Suppose three faces of σ are Möbius faces, and the remaining face is either a Möbius face or a triangle face. Then $M = L(4)$ or $L(5)$.*

Proof Suppose there are three Möbius faces and one triangle face in σ . First suppose that each edge of the triangle face is the core edge of a Möbius face. Then there is (up to symmetry) only one possibility, and one sees that the vertex link cannot be a sphere. Hence some edge of the triangle face is the boundary edge of a Möbius face. If the core edge of this Möbius face is not identified with one of the other edges of the triangle face, then not both of the remaining faces can be Möbius

faces. Hence it is identified with one of the other edges of the triangle face. Up to symmetry, there is again only one possibility and one finds a spine for $L(5)$ made up of one *twisted square* (see §3.1 for a definition) and one face.

Hence suppose there are four Möbius faces in σ . Given two Möbius faces, they meet along some edge e simply as a result of being faces of σ . Suppose e is a boundary edge of both Möbius faces. Then the condition that the other two faces are also Möbius forces the core edges to be identified such that a square is a spine for $L(4)$. Enumerating the remaining cases, one finds only one possibility up to symmetry, and this gives a spine for $L(5)$. ■

The following result follows from [4] (Proposition 5.3 and Corollary 5.4) and [5] (Lemma 7); see also the proof of Proposition 13 below.

Lemma 6 *Suppose M is a closed, orientable, irreducible 3-manifold with an arbitrary triangulation. If a face in the triangulation is a cone or a dunce hat, then the triangulation can be simplified to have fewer tetrahedra or M is homeomorphic with one of S^3 , $\mathbb{R}P^3$, or $L(3)$.*

The above observations lead us to the following definition:

Definition 7 (face-generic) *Suppose M is a closed, orientable 3-manifold. We say that a triangulation of M is face-generic if all faces are triangles or Möbius bands and each tetrahedron has at most two Möbius faces.*

Remark 8 *We have the following immediate consequences:*

- (1) *A face-generic triangulation may have more than one vertex. Examples are simplicial triangulations of any manifold, or the natural 2-vertex triangulations of lens spaces.*
- (2) *A face-generic triangulation has no edge of degree one, since such an edge can only be obtained from folding two faces of a tetrahedron together and hence creating at least one face that is a cone or dunce hat.*
- (3) *A 0-efficient triangulation of M is face-generic unless M is homeomorphic with one of S^3 , $L(3)$, $L(4)$ or $L(5)$.*
- (4) *A minimal triangulation of the closed, orientable, irreducible 3-manifold M is face-generic unless M is one of S^3 , $\mathbb{R}P^3$, $L(3)$, $L(4)$ or $L(5)$.*

Whilst it is easy to recognise whether or not a triangulation is face-generic, the class of face-generic triangulations is too large for many purposes. We now turn to pairs of faces to further restrict the class of triangulations we wish to consider.

2.4 Pairs of faces in triangulations

A combinatorial map will mean a continuous map between pseudo-manifolds $N \rightarrow M$ sending cells to cells and which restricted to the interior of each (possibly singular) simplex of each dimension is a homeomorphism onto its image. In particular, $\dim N \leq \dim M$. An example is the map $\tilde{\Delta} \rightarrow M$.

The following lemma follows from results by Burton [1, 2].




Figure 2: Identifications of two distinct faces along at least two edges: Picture (c) does not occur, and all other pictures allow an explicit modification of the triangulation making it smaller. The captions indicate which closed surfaces (if any) the map factors through.

Lemma 9 (Disc subcomplexes) *Let M be a closed, orientable, irreducible 3-manifold with a face-generic triangulation having at least 3 tetrahedra. Let D be a disc triangulated with 2 triangles meeting along two edges, and suppose that $f: D \rightarrow M$ is a combinatorial map with the property that the images of the two triangles in D are distinct faces in M .*

- (1) *If the two interior edges of D are mapped to distinct edges, then the triangulation can be modified to a triangulation of M having fewer tetrahedra.*
- (2) *If the two interior edges of D are mapped to the same edge, then the triangles in D are mapped to Möbius faces with distinct boundary edges, or the two faces form a spine for $L(4, 1)$ and hence $M = L(4, 1)$, or the triangulation can be modified to a triangulation of M having fewer tetrahedra.*

Proof The first part follows directly from Lemma 2.7 and Corollary 2.10 in [1], and Lemma 3.6 and Corollary 3.8 in [2].

For the second part, observe that if the two interior edges are identified, then, because we cannot get cones, dunce hats or 3-folds, the images are Möbius faces sharing at least the core edges. Suppose the two boundary edges are identified. There are two cases, depending on the orientation of the edges. One case gives a spine for $L(4, 1)$. In the other case, the map has image equivalent to Figure 2(f) and hence there is a map $f': D \rightarrow M$ as in part (1). ■

Remark 10 *The techniques of the lemma also apply to reducible manifolds. In this case, the modification procedure either results in a triangulation of M having fewer tetrahedra, or $M = M_1 \# M_2$, and one obtains triangulations of M_1 and M_2 , each having fewer tetrahedra.*

The above lemma leads us to the following definition (see also Figure 3):

Definition 11 (face-pair-reduced) *Let D be a disc triangulated with 2 triangles meeting along two edges. We will call a triangulation face-pair-reduced if every combinatorial map $f: D \rightarrow M$, which maps the two triangles in D to distinct faces in M , has the property that the two interior edges of D are mapped to the same edge, and the triangles in D are mapped to Möbius faces with distinct boundary edges.*

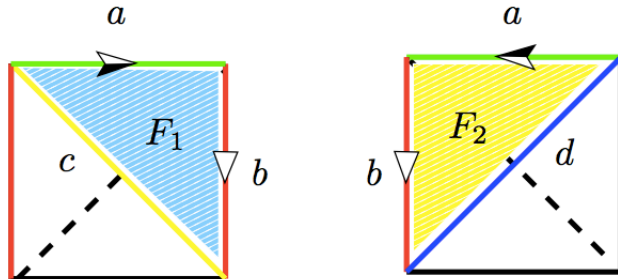


Figure 3: Face-pair-reduced. If $F_1 \neq F_2$, then $a = b$ and $c \neq d$.

Remark 12 *We have the following immediate consequences:*

- (1) *A tautological, but most useful, consequence of the definition is that case (1) in Lemma 9 cannot occur in a face-generic, face-pair-reduced triangulation. Thus, each face is uniquely characterised by two consecutive distinct, oriented edges in its boundary.*
- (2) *A face-pair-reduced triangulation may have more than one vertex. Examples are simplicial triangulations of any manifold, and the natural 2-vertex triangulations of lens spaces.*
- (3) *A face-pair-reduced triangulation has no cone faces (see Burton [1], Lemma 2.8).*
- (4) *A minimal triangulation of a closed, orientable, irreducible 3-manifold with at least 3 tetrahedra is face-pair-reduced.*
- (5) *A 0-efficient and face-generic triangulation may not be face-pair-reduced. An example is given by the 4 tetrahedron triangulation of S^3 defined by the following face pairings:*

Tetrahedron	Face 012	Face 013	Face 023	Face 123
0	2 (231)	1 (321)	0 (312)	0 (230)
1	3 (231)	3 (023)	2 (032)	0 (310)
2	3 (013)	3 (120)	1 (032)	0 (201)
3	2 (301)	2 (012)	1 (013)	1 (201)

This triangulation can be simplified to a 1-tetrahedron triangulation of S^3 via a sequence of three Pachner moves.

2.5 Obtaining small face-pair-reduced, face-generic triangulations

Every simplicial triangulation is face-pair-reduced and face-generic. In particular, one can convert any triangulation to a face-pair-reduced, face-generic triangulation by performing at most two barycentric subdivisions. However, this increases the number of tetrahedra.

Proposition 13 *There is an algorithm which takes as input any triangulation of a closed, orientable, irreducible 3-manifold M , and either outputs a face-pair-reduced, face-generic triangulation of M having at most the same number of tetrahedra and precisely one vertex, or concludes that M is S^3 , $\mathbb{R}P^3$, $L(3)$, $L(4)$ or $L(5)$.*

Proof For the purpose of this paper, we are merely interested in the existence of an algorithm, not in producing an efficient algorithm.

- (0) First run the 3–sphere recognition algorithm to decide whether M is S^3 . If yes, output M is S^3 . Else go to step (1).
- (1) If the triangulation has only one vertex, go to (2). Otherwise the given triangulation has more than one vertex, and one converts it to a 1–vertex triangulation, thereby reducing the number of tetrahedra (see the proofs of Theorem 5.5 and Proposition 5.1 in [4]). Then go to (2).
- (2) If there are at most 2 tetrahedra, one recognises the manifold and if one exists, outputs a face-pair-reduced, face-generic, 1–vertex triangulation of M , or outputs the homeomorphism type of M . Otherwise go to step (3).
- (3) If a face is a 3–fold, then output M is $L(3)$; else if some tetrahedron has more than 2 Möbius faces, then $M = L(4)$ or $L(5)$, and the conclusion can be obtained either from the combinatorics of the 2–skeleton, or by computing $H_1(M)$ from the 2–skeleton; else go to step (4).
- (4) Check if the triangulation is face-pair-reduced (recall that existence of cone or dunce faces implies not face-pair-reduced). If “Yes,” then the triangulation is face-pair-reduced and face-generic, and output the triangulation. If “No,” then apply the explicit operation from the proof of Lemma 2.7 in [1]. This either gives a triangulation of M having the same number of tetrahedra but more vertices, or it gives triangulations of M and S^3 , each having fewer tetrahedra. In the second case, one needs to run the 3–sphere recognition algorithm to determine which triangulation corresponds to M . In either case, one feeds the new triangulation of M back to (1).

In steps (2) and (3), the number of tetrahedra is not affected, so if the algorithm does not terminate in these steps, one arrives at (4) with the same number of tetrahedra. In step (4), the number of tetrahedra either decreases and one goes back to (1), or the number of tetrahedra remains the same, but subsequently decreases in step (1). In particular, if the algorithm does not terminate in steps (2), (3) or (4), then the loop (4) \rightarrow (1) \rightarrow (2) \rightarrow (3) \rightarrow (4) decreases the number of tetrahedra. It follows that the algorithm will terminate. ■

Remark 14 *For a reducible 3–manifold, M , the above algorithm can be adapted to produce face-pair-reduced, face-generic triangulations of (not necessarily prime) summands of M .*

3 Squares and tetrahedra in face-generic triangulations

The previous section focussed on the study of triangulations using faces and certain pairs of faces. We will now take the combinatorial study of triangulations further by examining twisted squares and pairs of twisted squares. This leads to a classification of the possible types of tetrahedra in a face-pair-reduced, face-generic triangulation.

3.1 Squares in triangulations

Let σ be an oriented tetrahedron in $\tilde{\Delta}$. A *twisted square* in σ is a properly embedded disc in σ such that the boundary of the disc is the union of two pairs of opposite edges of σ ; see Figure 4.

Twisted squares first appeared in work of Thurston [11]. Notice that each of the three quadrilateral types in σ naturally corresponds to the twisted square, which the quadrilateral meets in its corners. The adjective *twisted* will often be omitted.



Figure 4: Twisted square. On the left is a twisted square as realised by a fundamental piece of Schwarz's minimal surface, on the right is a quadrilateral disc and the boundary of its dual square.

A twisted square has the following geometric realisation. Identify σ with a regular Euclidean tetrahedron. Then the minimal disk spanning the union of two pairs of opposite edges of σ is a twisted square. This is a fundamental piece of Schwarz's minimal surface. We will work with twisted squares combinatorially, but alternative arguments using the Gauß–Bonnet formula can be given because this realisation of a twisted square is isometric with an equilateral hyperbolic polygon with all dihedral angles $\pi/3$.

Under the map $\tilde{\Delta} \rightarrow M$, the edges of a square may become identified; the image may have 4, 3, 2 or 1 distinct edges in M . Labelling the boundary of a square by letters indicating the *unoriented* edges in M , we obtain the following basic *partition-types of squares*:

$$[A] \ abcd \quad [B] \ abac \quad [C] \ aabc \quad [D] \ abab \quad [E] \ aaab \quad [F] \ aabb \quad [G] \ aaaa,$$

where different letters correspond to distinct (unoriented) edges in M .




Figure 5: The partition types of squares (edges of the same colour are identified; edges of distinct colours are distinct)

It will be useful to have names for certain types of squares, as we did for triangles. We will only need names indicating the topology. A square of type B can be an *annulus square* or a *Möbius*

square. A square of type C is a *Möbius square*, a square of type D is a *torus square* or *Klein square* or *projective square*, a square of type F is a *Klein square*. See Figure 6.

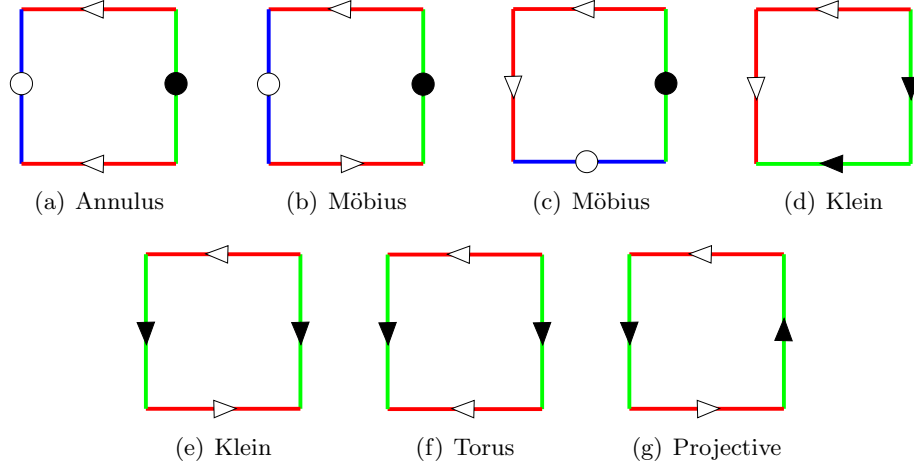


Figure 6: Some topological types of squares in a face generic triangulation

We now define what we mean by a *pinched surface* (see also Figure 7). Suppose S is a compact surface and partition S into finite subsets with the property that all but finitely many of these sets are singletons. Then the quotient space is a *pinched surface*. Equivalently, consider a finite collection of pairwise disjoint subsets V_k of S . Then identifying all points in V_k to one point for each k gives a pinched surface.




Figure 7: A pinched surface

Lemma 15 *Suppose M is closed, orientable and has a face-generic triangulation.*

- (1) *No square is of type G .*
- (2) *If there is a square of type F , then M contains an embedded Klein bottle.*
- (3) *If there is a square of type D , and this square is not an embedded torus in M , then M either contains an embedded projective plane or an embedded Klein bottle.*
- (4) *If the triangulation is also face-pair-reduced, then no square of type D is a projective square.*

Proof Since the triangulation is face-generic, no face is a cone, 3-fold or dunce hat. If two consecutive edges of a square are identified, then the face containing them is a Möbius face. So if some square is of type G , then some tetrahedron has four Möbius faces. This completes the first case. The second also follows from the induced orientations of the edges. The third follows from the

observation that the possible orientations on the edges of a square of type D either give an embedded torus, and embedded Klein bottle or a pinched projective plane. Intersecting the pinched projective plane with the boundary of a small neighbourhood of the vertex gives two disjoint simple closed curves, and capping these off with disjoint discs in the vertex neighbourhood gives an embedded projective plane in M .

The fourth claim follows by drawing in the remaining edges of a tetrahedron containing a square of type D . The two faces are forced to be identified due to Definition 11. However, the induced face pairing will be orientation preserving, contradicting the hypothesis that M is orientable. ■

3.2 Pinched, capped and surgery 2–square surfaces

In this section, we describe certain unions of two twisted squares in distinct tetrahedra which allow us to draw conclusions about the topology of a triangulated 3–manifold.

Lemma 16 *Let M be a closed, orientable, triangulated 3–manifold, and S be a closed, connected surface which has a CW decomposition, whose set of 2–cells consists of at most two squares. Then $-1 \leq \chi(S) \leq 2$. Suppose further that $f: S \rightarrow M$ is a combinatorial map that takes the squares to twisted squares in distinct tetrahedra, and distinct edges in S to distinct edges in M . Then there is an embedding $g: S \rightarrow M$.*

Proof Since there are four edges, two faces and at least one vertex, we have $-1 \leq \chi(S) \leq 2$. In fact, S can be viewed as obtained from a hexagon by identifying boundary edges in pairs. Now $f: S \rightarrow M$ is a combinatorial map that takes the squares to squares in distinct tetrahedra, and distinct edges to distinct edges. Then $f(S)$ is a (possibly) pinched surface in M . We can replace $f(S)$ by an embedded surface homeomorphic with S as follows. The surface $f(S)$ is embedded outside the vertex neighbourhoods. The intersection with sufficiently small vertex linking spheres consists of a finite number of simple closed loops, which we may cap off using discs (starting from innermost circles) to obtain an embedded surface homeomorphic with S . This surgery of f gives the desired embedding $g: S \rightarrow M$. ■

We call the map $f: S \rightarrow M$ in the statement of Lemma 16 a *pinched 2–square surface* in M , and the embedding $g: S \rightarrow M$ of the proof a *capped 2–square surface*.

Now we take this analysis further in order to detect interesting surfaces in M even if there are more edge identifications amongst two twisted squares. We are not trying to get a complete list, but rather list the situations that arise in this paper, and call the resulting surfaces *surgery 2–square surfaces*. The basic fact we use is that if a simple closed curve on the boundary of a solid torus runs along the longitude $2n$ times, then it bounds a properly embedded surface of non-orientable genus n with one boundary component in the solid torus. For instance, if $n = 2$ then the surface is $(P\#P) \setminus \text{disc}$. We call such a curve a $2n$ –curve on the solid torus, and we include the case $n = 0$ to give a (boundary parallel or meridian) disc.

To illustrate the basic technique, suppose a square of type G in M has its boundary edge oriented coherently. Then the edge is a loop in M , and intersecting the square with a regular tubular neighbourhood of that edge in M gives a 4–curve on the neighbourhood, and hence bounds a $(P\#P) \setminus (\text{disc})$ in the tubular neighbourhood. Moreover, we can cap this off with the complement in the square of the neighbourhood and hence get an embedded Klein bottle in M . One can view

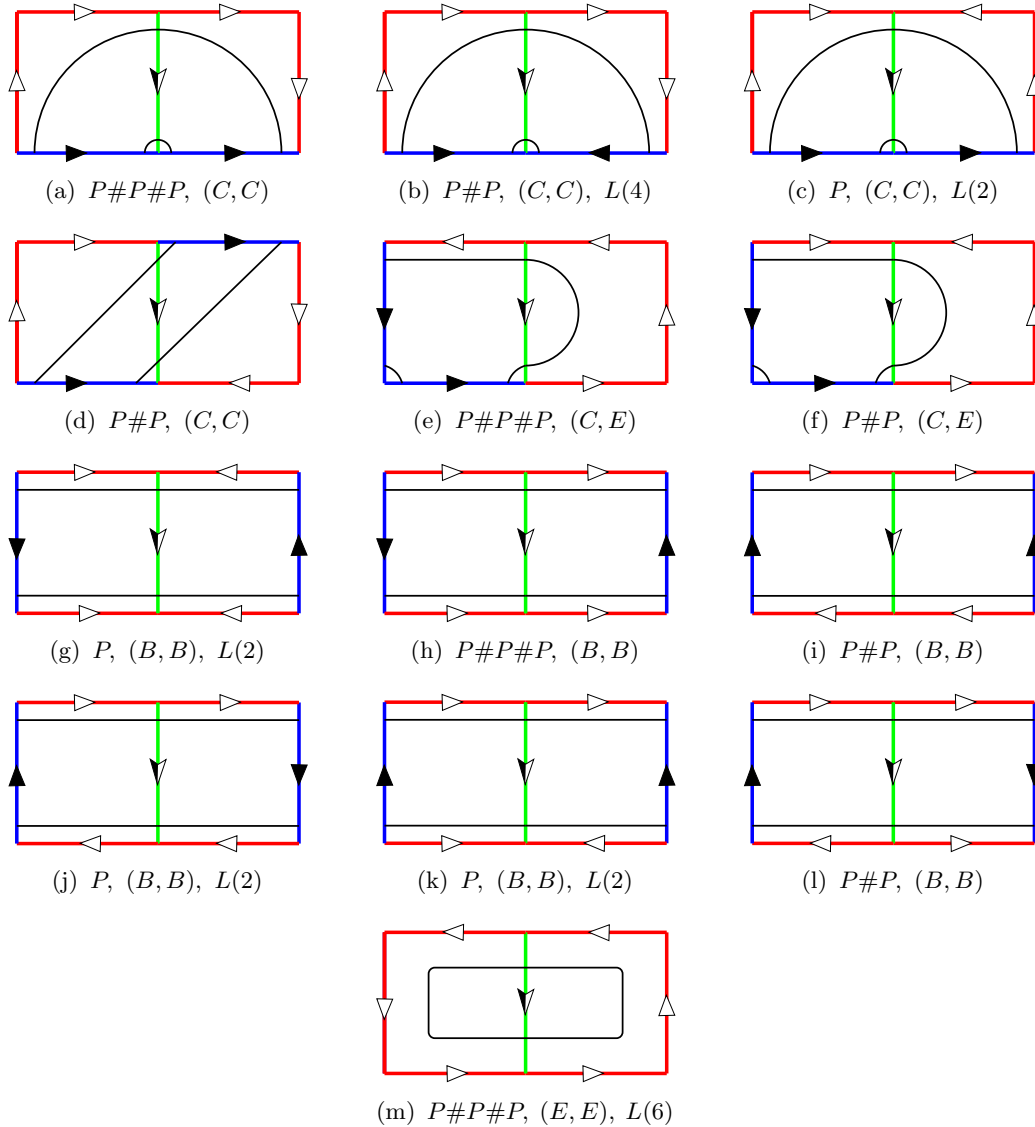


Figure 8: Pairs of squares in distinct tetrahedra detecting non-orientable surfaces or lens spaces (or lens space summands in case the manifold is reducible). Edges of distinct colours are distinct. The Klein bottles in 8(d) and 8(i) consist of an annulus in the edge neighbourhood and the annulus on the squares. Notice that some of these pictures are equivalent once one deletes the green diagonal from the hexagon.

the square of type G as the image of an immersion of a 1–square Klein bottle, and the surgery procedure has recovered an embedding. In practice, reference to a domain is not required, as the object given to us is the union of two twisted squares in M .

Some constructions with two squares are shown in Figure 8, and they should be interpreted as follows. Let N denote a small regular neighbourhood of the red edge. Suppose $X \subset M$ is the union of two twisted squares contained in distinct tetrahedra. Then X meets ∂N in the shown black curve(s). The intersection $N \cap X$ is replaced by an embedded surface in N , which has the same boundary curve(s). In almost all cases, $X \cap \partial N$ is connected, and (after choosing an orientation) its class in $H_1(N)$ can be read off from the picture, taking orientations into account. In some cases, $X \cap \partial N$ has one component which is zero in $H_1(N)$ (and is hence capped off with a disc) and another which may or may not be trivial. In Figures 8(d) and 8(i), one has two parallel non-trivial curves and these are connected by an annulus in N . In this case, an annulus on X gives an orientation reversing isotopy between these curves; whence the surface is a Klein bottle. We will call a surface that arises as described above a *surgery 2–square surface* in M .

3.3 Recognising projective planes using three squares

We continue the analysis of the previous section with three situations that arise at the end of this paper. Suppose there are three pairwise distinct tetrahedra σ_k containing squares $S_k \subset \sigma_k$ such that the edges are identified as shown in Figure 9. In the left-most picture, the subcomplex X formed by these squares has an edge e of degree 6 (shown in red). The intersection of the boundary of a small regular neighbourhood N of e with X results in a curve, which is embedded in ∂N and is null-homotopic in N . It hence bounds a disc in N . This curve also bounds a Möbius band in X , whence M contains an embedded projective plane. A similar surgery argument results in a projective plane in the middle picture, whilst in the last picture, merely capping is required to yield such a surface.




Figure 9: Three squares contained in three pairwise distinct tetrahedra giving a projective plane

3.4 Partition-types of tetrahedra

We now describe the combinatorics of a tetrahedron in a face-generic 1–vertex triangulation using the three squares contained in it. The tetrahedron σ in M is said to be of *partition-type* XYZ if the three squares supported by it are of partition-types X , Y and Z respectively. The labels are ordered lexicographically.

Lemma 17 Suppose M is closed, orientable and has a face-generic triangulation. Suppose σ is a tetrahedron in M , and let n be the number of pairwise distinct edges of σ . Then either

- (1) $n = 6$ and σ is of partition-type AAA ; or
- (2) $n = 5$ and σ is of partition-type AAC or ABB ; or
- (3) $n = 4$ and σ is of partition-type AAF , ABE , ACC , BBC or BBD ; or
- (4) $n = 3$ and σ is of partition-type BBF , BDE or DDD .

Moreover, in case DDD , all three squares in σ are embedded Klein bottles in M , unless one square gives a pinched projective plane and hence an embedded projective plane in M .

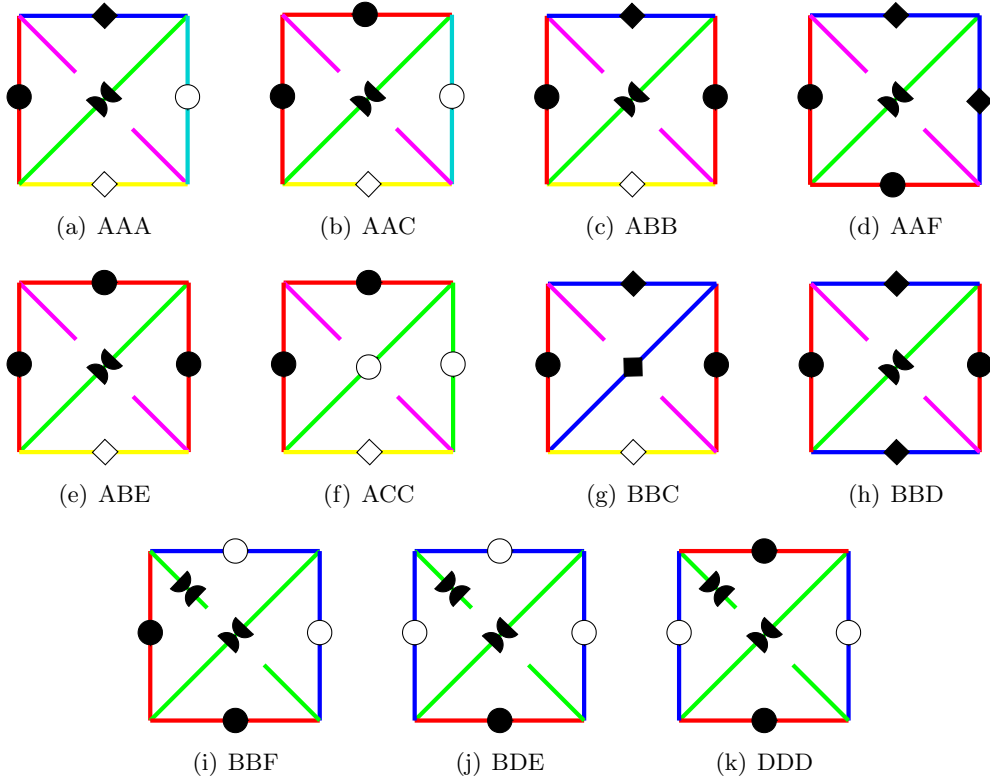


Figure 10: The partition types of tetrahedra (edges of the same colour are identified; edges of distinct colours are distinct; faces may or may not be identified)

Proof Recall that face-generic implies that all faces are Möbius faces or triangles, and σ has at most two Möbius faces. In particular, no square can be of type G .

Next suppose that some square is of type F . Since no face is a cone, the two remaining edges of σ are distinct from the edges of the square. If these two edges are identified, we have type BBF and $n = 3$, and if they are distinct, we have type AAF and $n = 4$.

Hence suppose that no square is of type F or G . We enumerate the remaining cases depending on n . First suppose $n = 6$. Then each square must be of partition type A .

Suppose $n = 5$. It follows that either a pair of opposite edges is identified, giving case ABB , or a pair of edges incident with a face is identified, giving AAC .

Suppose $n = 4$. First assume that three edges are identified. Three given edges are either incident with one vertex, or contained in a unique face or in a unique square. The three edges cannot be incident with one vertex, since then there would be a cone face. Also, the three edges of a face cannot be identified since otherwise there would be a 3-fold face or a dunce face. Hence one of the squares is of type E . Since this square contains exactly two edges of the triangulation of M , the remaining two edges must be distinct and distinct from the edges in E , giving ABE . Next assume that two pairs of edges are identified. If two pairs of opposite edges are identified, we obtain BBD . If one pair of opposite edges is identified and one pair of adjacent edges, then we obtain BBC . If no pair of opposite edges is identified (and hence two pairs of adjacent edges), we obtain ACC .

Suppose $n = 3$. No square can be of type A . Suppose there is a square of type B . Then it contains the three distinct edges of σ . Enumerating the possible identifications of the remaining two edges with the three edges of the B -square gives BDE , BBF , BBG , BCE . We have already assumed that no square is of type F or G , and in case BCE , there would be 3 Möbius faces. Hence the only possibility is BDE . Now then suppose no square is of type B , but there is a square of type C . In each possible case, one finds three or four Möbius faces. Hence suppose no squares are of type B or C . If there is a square of type D , then one of the remaining edges must be different, and since we don't have a type B square, the remaining edges are identified. This gives DDD . Now one of the squares gives a projective plane unless all three squares are Klein bottles. Last, suppose all squares are of type E . This is not possible.

Suppose $n = 2$. In this case, we can only have squares of type D or E . The only possibility is DEE , which gives four Möbius faces. ■

Corollary 18 *Suppose M is closed, orientable and has a face-generic triangulation without Klein or projective squares. Then each tetrahedron is combinatorially equivalent to one of eight possibilities; one each falling into the partition types AAA , AAC , ABB , ABE , ACC , BBC , $sBBD$, $sBDE$ as shown in Figure 11. In the first six cases, this is just the generic type (up to combinatorial equivalence), but in the latter two, there are prescribed edge identifications.*

If, in addition, the triangulation is face-pair-reduced, then

- (1) *case BBC does not occur, and*
- (2) *the two Möbius faces of every tetrahedron of type $sBDE$ in M are identified (i.e. such a tetrahedron is a so-called 1-tetrahedron layered solid torus $LST(3, 2, 1)$).*

Proof Refer to the types listed in Lemma 17. In each of the cases AAF , BBF , and DDD , there is a Klein or projective square, and hence they cannot occur.

Since we assume there are no Klein or projective squares, the squares of type D in BBD and BDE give tori. This determines $sBBD$, and also, together with the hypothesis that there are no cone faces, determines $sBDE$. ■

4 Extremal rays with support in one tetrahedron

There is an algebraic approach to normal surface theory due to Haken (see [4]), which arises from the fact that the normal discs in the cell structure of a normal surface match up across faces in the triangulation. Another algebraic approach (which goes back to Thurston and Jaco and can be found

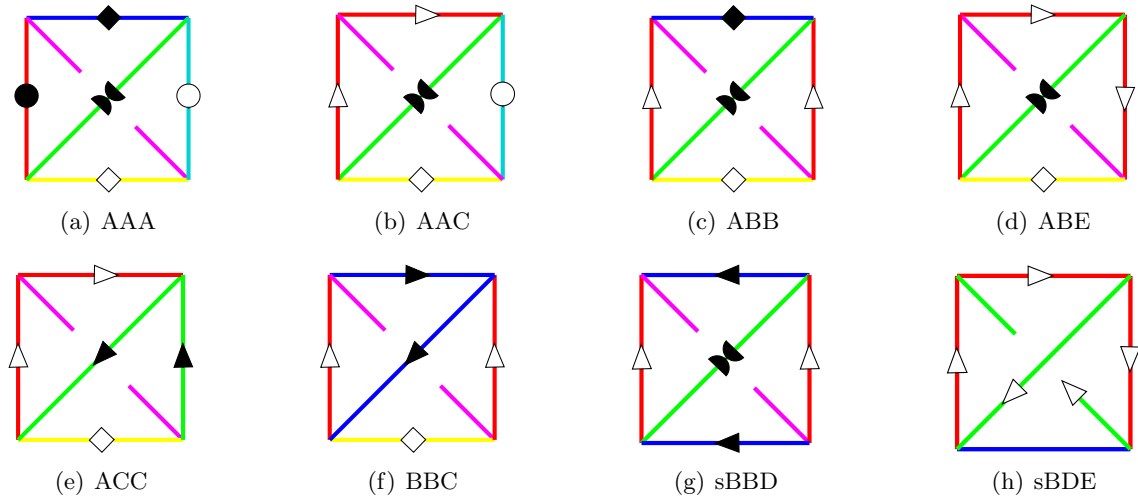


Figure 11: The remaining possibilities for tetrahedra up to combinatorial equivalence (edges of the same colour are identified; edges of distinct colours are distinct)

in Tollefson [13]), relies on the key observation that each normal surface is uniquely determined (up to normal isotopy) by the quadrilateral discs in its cell structure, and that the normal surface meets a small neighbourhood of each edge of the triangulation in a disjoint union of discs. Intuitively, these equations arise from the fact that as one circumnavigates the earth, one crosses the equator from north to south as often as one crosses it from south to north. In terms of the abstract neighbourhood of an edge, one picks an orientation of the edge (so that north and south are defined) and observes that triangles meeting the edge remain in one of the hemispheres, whilst quadrilaterals either cross from north to south (in which case the corner of the quadrilateral at e is given sign -1) or from south to north (in which case the corner of the quadrilateral at e is given sign $+1$). See Figure 12.




Figure 12: Slopes of quadrilaterals

Given the edge e and a quadrilateral type q , the total weight $wt_e(q)$ of q at e is the sum of signs over all corners of q at e (where the empty sum has value 0). For the quadrilateral type q introduce the variable $x(q)$. The Q -matching equation at e is then defined as $0 = \sum_q wt_e(q)x(q)$, where the sum is taken over all quadrilateral types. For instance, placing a quadrilateral of each type in a fixed tetrahedron σ gives a so-called *tetrahedral solution* to these equations: this has $x(q) = 1$ for each quadrilateral type q supported by σ and $x(q) = 0$ for all others. See [13, 12] for more details.

We call a solution to the matching or Q -matching equations a *1-quad type solution* if it has exactly one non-zero quadrilateral coordinate, and a *2-quad type solution* if it has exactly two non-zero quadrilateral coordinates. The non-zero quadrilateral coordinates need not have the same sign. If x is *any* solution to the matching or Q -matching equations with support in exactly one tetrahedron, then one may add a multiple of a tetrahedral solution to x so that the result has at most two quadrilateral coordinates that are non-zero and so that both are non-negative.

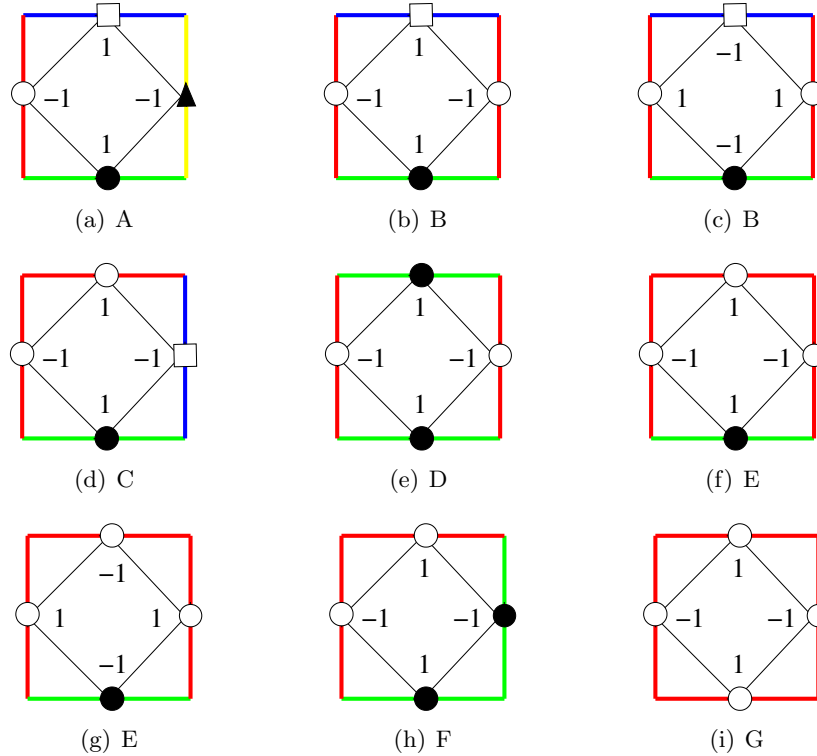


Figure 13: The signs of corners of quadrilateral discs shown with the dual twisted squares (edges of the same colour are identified; edges of distinct colours are distinct)

Lemma 19 *Let M be a closed, orientable 3-manifold with an arbitrary triangulation. The tetrahedron σ in M supports a 1-quad type solution if and only if at least one square in σ is of type F or G .*

If the triangulation is face-generic, then M supports a 1-quad type solution if and only if it contains a Klein square.

If the triangulation is 0-efficient and no face is a cone, then M supports a 1-quad type solution if and only if it either contains a Klein square or a square is a spine for $L(4, 1)$.

Proof In an orientable 3-manifold, the signs at the corners of a quadrilateral disc in a tetrahedron are the same at opposite corners and opposite at adjacent corners (see [12]). A quadrilateral disc has its four corners on the boundary of a twisted square; see Figure 13. There is a 1-quad type solution if and only if the signs of the square at each edge sum to zero. It can now be deduced from the figure that there is a 1-quad type solution if and only if there is a twisted square of type F or G . This proves the first assertion.

The second assertion follows as in the proof of Lemma 15.

Hence assume that the triangulation is 0-efficient and no face is a cone. Then a square of type F is a Klein square and a square of type G is a spine for $L(4, 1)$. ■

The lemma implies that the existence of a 1-quad type solution is easy to deduce from the combinatorics of a triangulation.

Proposition 20 *Suppose a triangulation of the closed, orientable 3-manifold M supports no 1-quad type solution.*

If the triangulation is face-generic, then there is no 2-quad type solution supported by precisely one tetrahedron σ .

If the triangulation is 0-efficient and there is a 2-quad type solution supported by precisely one tetrahedron σ , then either a face of σ is a cone (in which case M is homeomorphic with S^3) or σ is of type DEE (in which case M is homeomorphic with $L(5)$).

Proof Suppose the 2-quad type solution x does not vanish on the two quadrilateral coordinates $p, q < \sigma$. If $x(p) = x(q)$, then one can adjust by a tetrahedral solution to obtain a 1-quad type solution which does not vanish precisely on the remaining quadrilateral coordinate supported by σ , contradicting the assumption that there is no 1-quad type solution.

After possibly adjusting by a tetrahedral solution and relabelling, we may assume that $x(p) > x(q) > 0$. Let $\alpha = x(p)$ and $\beta = x(q)$. Since the space of all solutions to the Q -matching equations is a rational polyhedral cone and there is no 1-quad type solution, it follows that α and β are linearly dependent over \mathbb{Q} . This fact, however, is irrelevant.

Up to orientation, we may assume that the weights on the six edges of σ are $\alpha, \alpha, \beta - \alpha, \beta - \alpha, -\beta, -\beta$. Since there are precisely two positive weights, and we assume that there is no 1-quad type solution, it follows that the two edges of σ with weight α are distinct in M . Since $\alpha \neq \beta$, it follows that each edge of positive weight is identified with at least two edges of negative weight. Whence with precisely two edges of negative weight. Up to symmetry and ignoring orientation of edges, this gives precisely two possible patterns for the identification of the edges of σ . In one case, there are three edges incident with a vertex of σ , which must be identified. But then a face is a cone, giving a contradiction when the triangulation is face-generic, and the conclusion that M is S^3 when the triangulation is 0-efficient.

In the other case, σ is of type DEE and a square and a face of σ form a spine for $L(5)$. However, there is no tetrahedron of type DEE in a face-generic triangulation according to Corollary 18. ■

Corollary 21 *If there is a 1-quad type or 2-quad type solution supported by a single tetrahedron in a 0-efficient triangulation, then the underlying manifold is S^3 , $L(5)$ or contains an embedded Klein bottle. Moreover, a face-generic triangulation without Klein squares does not admit such a solution.*

As an example, the Seifert fibered space $S^2((2, 1), (2, 1), (3, -1))$ has a minimal 4-tetrahedron triangulations with a 1-quad type solution; this is triangulation #1 of this manifold in Burton's closed, orientable manifold census [3]. This triangulation is 0-efficient, face-generic and contains a Klein square.

Corollary 22 *Let M be a closed, orientable 3-manifold with face-generic triangulation \mathcal{T} . If there are no Klein squares, then each tetrahedral solution lies on an extremal ray in $Q(\mathcal{T})$, the polyhedral cone of all non-negative solutions to the Q -matching equations.*

5 Clusters of three 2–quad type solutions

We are now in a position to state, and prove, the main application of our study of face-pair-reduced, face-generic triangulations:

Theorem 23 *Suppose M is a closed, orientable 3–manifold with face-pair-reduced, face-generic triangulation such that*

- (1) *there are no Klein squares,*
- (2) *no two squares are identified as shown in Figure 8,*
- (3) *no three squares are identified as shown in Figure 9,*
- (4) *no capped 2–square surface is a non-separating torus,*
- (5) *no capped 2–square surface is non-orientable.*

Then there is no cluster of three 2–quad type solutions.

Remark 24 *The two minimal triangulations of $\mathbb{R}P^3$ have clusters of three 2–quad type solutions. However, one of these triangulations is face-generic and not face-pair-reduced, and the other is not face-generic.*

Proof of Theorem 1 using Theorem 23 Let M be a closed, orientable 3–manifold. We may suppose that M is irreducible and atoroidal (since otherwise we are in the first or second case) and that we are given a minimal triangulation. According to §§2.3–2.4, either the triangulation is face-pair-reduced and face-generic or M is a Seifert fibered space (which yields the third case). Hence assume that the minimal triangulation is face-pair-reduced, face-generic and that we have a cluster of three 2–quad type solutions. Then one of the five hypotheses in Theorem 23 does not hold.

The negation of (1) implies that M is toroidal or a Seifert fibered space.

The negation of (2) or (5) implies that M contains an embedded non-orientable surface of Euler characteristic -1 , 0 or 1 . In the first case, there is an embedded non-orientable surface of genus 3, in the second M is toroidal or Seifert fibered and in the third M is real projective 3–space and hence Seifert fibered.

The negation of (3) implies that M is real projective 3–space.

The negation of (4) is that M contains a non-separating torus. This is homologically non-trivial. Since M is irreducible, every sphere is homologically trivial, and hence the torus must be incompressible, and hence M is toroidal. This concludes the proof of Theorem 1. ■

Proof of Theorem 23 For the remainder of this section, suppose the hypotheses of Theorem 23 are satisfied and that the triangulation has at least 3 tetrahedra, but that there is a cluster of three 2–quad type solutions. Suppose that M is oriented and that all tetrahedra are oriented coherently. Denote the solutions x_1, x_2, x_3 and suppose σ_0 is a tetrahedron in their common support. We denote the pairwise distinct quadrilateral types $q_1, q_2, q_3 < \sigma_0$ such that $x_i(q_i) \neq 0$. For each solution x_i , there also is tetrahedron σ_i and quadrilateral $p_i < \sigma_i$ such that $x_i(p_i) \neq 0$. The tetrahedra $\sigma_1, \sigma_2, \sigma_3$ are not necessarily pairwise distinct. However, since the triangulation is face-generic and has no Klein squares, it follows from Corollary 21 that σ_1, σ_2 and σ_3 are distinct from σ_0 .

Outline: A classification of the possible types of tetrahedra in a face-pair-reduced, face-generic triangulation was given in Corollary 18. We will first focus on one 2-quad type solution x supported on quadrilateral types q and p and determine the possible partition-types for the squares associated with the quadrilateral types. With this information, we then show that σ_0 has no square of type B . We finally show that σ_0 has no square of type A , hence contradicting the classification of tetrahedra in Corollary 18.

Orientation: In some arguments, we will only analyse certain situations up to combinatorial equivalence or up to orientation in order to limit the number of cases to be considered. However, this has to be taken into account when more than one of these situations are put together into a common framework since M is oriented and all tetrahedra are oriented coherently.

Partition-types of quadrilaterals: The Q -matching equations and signs of corners of quadrilateral discs have been reviewed in §4. Let σ be an oriented tetrahedron in M , containing the quadrilateral disc q . Then q has signs $+1, +1, -1, -1$ at the edges of a twisted square in σ . In M , these signs could be at 4, 3, 2 or 1 distinct edges as shown in Figure 13. We list these cases according to the partition type of the square containing the corners of the quadrilateral, and refer to this as the *partition type* of the quadrilateral.

The notation we use is explained with an example. The partition $(1, -1 \mid 1 \mid -1)$ indicates that the signs are at three distinct edges in M and the respective weights are 1 and -1 at one edge, 1 at the second edge and -1 at the third edge. Lemma 19 implies that no square is of type F or G since the triangulation is face generic and has no Klein square, so we have the following cases:

- A $(1 \mid 1 \mid -1 \mid -1)$
- B $(1, 1 \mid -1 \mid -1)$ or $(-1, -1 \mid 1 \mid 1)$
- C $(1, -1 \mid 1 \mid -1)$
- D $(-1, -1 \mid 1, 1)$
- E $(1, 1, -1 \mid -1)$ or $(-1, -1, 1 \mid 1)$

Let x be a 2-quad type solution to the Q -matching equations with support in two distinct tetrahedra, σ_a and σ_b . After possibly scaling and relabelling the tetrahedra, we may assume that $x(q) = 1$ and $x(p) = t$, where $q < \sigma_a$, $p < \sigma_b$ and $0 < |t| \leq 1$. To simplify notation, we let $s = |t|$. Note that according to Lemma 19 and Proposition 20, neither σ_a nor σ_b supports a 1-quad type or 2-quad type solution, and $q < \sigma_a$ and $p < \sigma_b$ are of partition-type A, B, C, D or E . If q is of partition-type X and p is of partition-type Y , we will say that x is a partition-type (X, Y) solution. We will now determine the possible partition-types (X, Y) , as well as the corresponding weights $(1, s)$.

We first consider the possible weights on the edges of σ_b . It follows from the Q -matching equations, that the sum of all weights at each edge has to equal zero. Since the partition type takes the number of pairwise distinct edges and their weights into account, in order to determine which types can possibly balance other types, one needs to determine how many edges have non-zero weight. This is given by the *reduced partition types of quadrilaterals*, where only non-zero weight sums are counted:

- A' $(1 \mid 1 \mid -1 \mid -1)$
- B' $(2 \mid -1 \mid -1)$ or $(-2 \mid 1 \mid 1)$
- C' $(1 \mid -1)$
- D' $(-2 \mid 2)$
- E' $(1 \mid -1)$

Since no further identifications are allowed between the edges of a given square, it follows that the only potential partition-types (X, Y) for the solution x are

- (X, X) , where $X \in \{A, B, C, D, E\}$, or
- (X, Y) , where $X, Y \in \{C, D, E\}$ and $X \neq Y$.

We now narrow down the profile of *some* of the partition-types using the five hypotheses of Theorem 23. Our strategy is to consider the twisted squares $S_q \subset \sigma_a$ and $S_p \subset \sigma_b$ dual to $q < \sigma_a$ and $p < \sigma_b$ respectively.

Claim 0: For each of the partition types (A, A) , (B, B) , (C, C) , (C, D) , (D, C) there is (up to combinatorial equivalence) a unique identification between the edges of S_q and S_p as shown in Figure 14. Moreover, the type (C, E) does not occur.




Figure 14: Shown tetrahedra are coherently oriented in each pair (but not necessarily between different pairs); edges of distinct colours are distinct; except the black edges could be any colour (restrictions apply) and do not need to be of the same colour; the shaded faces are identified.

We will prove the claim one partition type at a time:

(A, A) : Each of the squares S_q and S_p has four pairwise distinct edges in M . The signs of the corners of q and p sum to zero at each edge in M , and hence each edge in S_q is identified with a unique edge of S_p and the signs at these edges are opposite. In particular, the union of S_q and S_p gives a capped 2-square surface S . Due to hypothesis (5), S must be orientable. Lemma 16 states that $-1 \leq \chi(S) \leq 2$, and hence S is a torus or sphere. Since the signs of corners of p and q cancel in pairs and alternate around the squares, there are only two possible pictures, a hexagonal torus and a sphere. The hexagonal torus is shown in Figure 14(a), and we will now show that the case of a sphere does not occur.

In this case, the identifications of the edges of the squares are as shown in Figure 15. We may assume that σ_a is the left-hand tetrahedron in each of Figure 15(a) and 15(b) and oriented as shown, and there are two possibilities for the orientation of σ_b . Definition 11 implies that the two faces shaded in Figure 15(a) must be identified, which is not possible in an oriented 3-manifold.

Applying Definition 11 to Figure 15(b) using the red and green edges forces the two shaded faces to be identified. The remaining "back faces" are similarly identified using the blue and yellow edges. The same reasoning applies to the front faces. Thus, the triangulation of M consists of exactly two tetrahedra (it is in fact a triangulation of S^3 with 4 vertices and 6 edges), contradicting the fact that we have at least 3 tetrahedra.




Figure 15: Shown tetrahedra for (A, A) are coherently oriented in each pair; edges of distinct colours are distinct, except the black edges could be any colour (restrictions apply) and do not need to be of the same colour.

(B, B) : If each of the squares S_q and S_p is of type B , then their reduced partition types must be $(2 \mid -1 \mid -1)$ and $(-2 \mid 1 \mid 1)$, with the edges of weights 2 and -2 identified. A square of type B is either an annulus or a Möbius square. This gives three main cases.

Suppose that one square is a Möbius square and the other is an annulus square. Up to combinatorial equivalence, this gives two possibilities, shown in Figures 8(k) and 8(l).

Next, suppose that both squares are Möbius squares. Up to combinatorial equivalence, this gives three possibilities. Two of these possibilities are shown in Figures 8(i) and 8(j). The remaining is shown in Figure 16(a). One now argues as in case (A, A) : there are two possible cases depending on the orientations of σ_a and σ_b . Definition 11 again implies that one of these cannot occur due to the orientability of M , and the other forces the four faces of σ_a and σ_b to be identified, giving a two tetrahedron triangulation. We already have the desired contradiction, but would like to point out that the four face pairings imply that all squares in the identification space are of type A , contradicting the fact that S_q and S_p are of type B . So this situation cannot occur in a face-generic, face-pair reduced, orientable 3-manifold triangulation.




Figure 16: Shown are the squares for the remaining possibilities for (B, B) , which cannot occur in a face-generic, face-pair reduced, orientable 3-manifold triangulation.

Last, suppose both squares are annulus squares. Up to combinatorial equivalence, this gives four

possibilities. Two of these possibilities are shown in Figures 8(g) and 8(h). The third is shown in Figure 16(b) and, exactly as above, ruled out with Definition 11. This only leaves the one possibility shown in Figure 14(b).

(C, C): Each square of type C is topologically a Möbius band, possibly with some vertices identified. Since a square of type C has reduced partition type $(1 \mid -1)$, each of the edges of degree one in one square is identified with a degree one edge of opposite sign of the other square. If this is the only identification between edges of the squares, then the union of S_q and S_p is a *non-orientable* pinched 2-square surface (either a Klein bottle or $P^2 \# P^2 \# P^2$). We therefore have a non-orientable capped 2-square surface in M , contradicting hypothesis (5).

The only other possibility is that, in addition, the degree two edges in S_q and S_p are also identified, so the 2-complex made up of S_q and S_p has one degree 4 edge and two degree 2 edges. Up to combinatorial equivalence, there are six possibilities. Four of these are shown in Figures 8(a)–8(d). The fifth possibility is shown in Figure 17, and, as above, Definition 11 implies that this is not possible. The last possibility is shown in Figure 14(c).

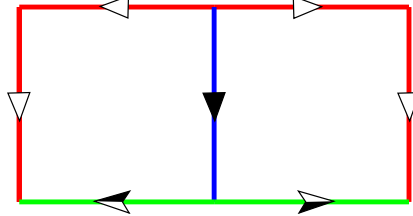


Figure 17: Shown are the squares for the fifth possibility for (C, C) , which cannot occur in a face-generic, face-pair reduced, orientable 3-manifold triangulation.

(C, D) or (D, C): Lemma 15 and hypothesis (1) imply that the square of type D is a torus square. So up to combinatorial equivalence, there is only one possibility for the identification of the two squares. Moreover, Definition 11 implies that two faces of σ_a and σ_b are identified and thus, up to orientation, there is only one possibility, shown in Figure 14(d). The face identification yields an identification of two *diagonals*. Definition 11 also implies that this diagonal cannot be identified with any of the edges in the square of type D , and face-generic implies that it can also not be identified with the core edge of the Möbius square of type C .

(C, E): This case does not occur: the only two possibilities are shown in Figures 8(e) and 8(f), giving either a Klein bottle or $P^2 \# P^2 \# P^2$.

Claim 1: The tetrahedron σ_0 is of type AAA , AAC or ACC .

To prove the claim, suppose that at least one of the 2-quad type solutions is of type B in σ_0 . According to *Claim 0*, the only possibility for a companion for a type B solution is type B and as shown in Figure 16(a). In particular, the squares of type B must be annuli. Now we assume that there are three 2-quad type solutions, and their types are given by the type of σ_0 . Refer to the types of tetrahedra in Figure 11. In ABB , BBC , $sBBD$, $sBDE$, one square of type B is a Möbius square. Hence σ_0 is of type ABE .

We know that a companion for a square of type A (respectively B) is also of type A (respectively B). Suppose that the companion square of type A is supported by σ_1 , and the companion square of type B is supported by σ_2 .

If $\sigma_1 \neq \sigma_2$, then (using the classification from *Claim 0*), the square of type A in σ_1 , the square of type B in σ_2 and the square of type E in σ_0 form a 2-complex that is combinatorially equivalent to the one shown in Figure 9(a), contradicting the hypothesis of Theorem 23.

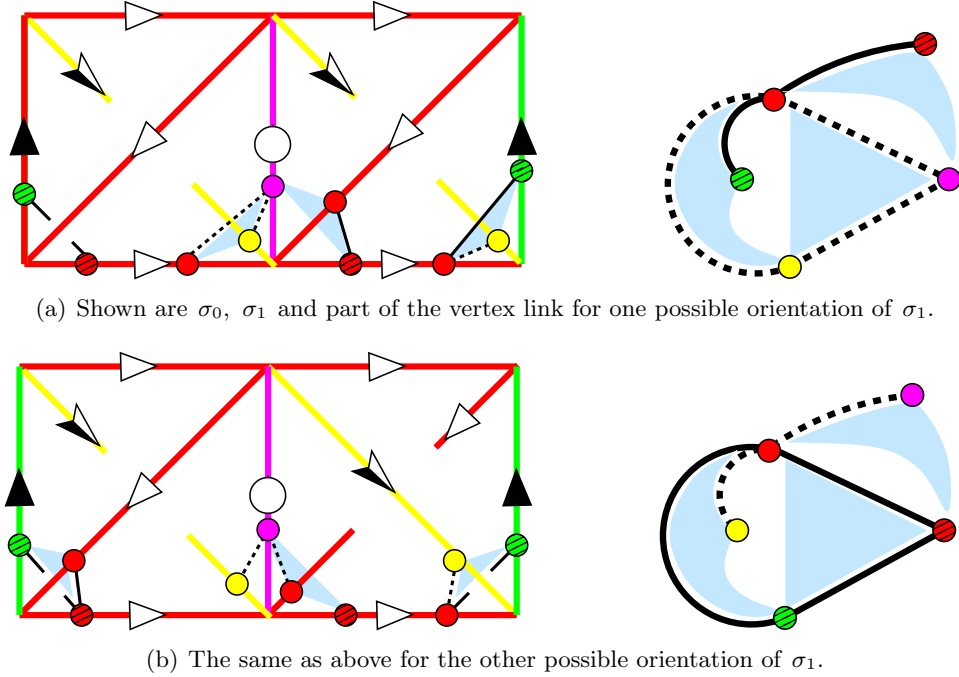


Figure 18: (ABE) The two possible orientations for the companion of a tetrahedron of type ABE. In each case, the indicated normal arcs form two simple closed curves that meet transversely in one point on the vertex link (the shaded triangles in the link are viewed from the vertex).

Hence $\sigma_1 = \sigma_2$. Without loss of generality, we may assume that the tetrahedron σ_0 of type ABE is oriented as in Figure 11. Since σ_1 contains squares of type A and B , this tetrahedron is either of type ABE or of type ABB , due to the classification in Figure 11. However, if σ_1 is of type ABB , then the two squares of type B in σ_1 meet in a single edge e in M , and this edge is not contained in the square of type A in σ_1 . However, in σ_0 this edge is in the intersection of the squares of types A and B . Whence σ_1 is of type ABE .

We need to distinguish two cases, depending on whether σ_1 is oriented as in Figure 11 or whether it is oppositely oriented. These cases are shown in Figure 18 (a) and (b) respectively. Examining the normal curves indicated in Figure 18 shows that the dashed arcs form a simple closed loop on the vertex link, and so do the solid arcs, but these two loops meet transversely in a single point. This contradicts the fact that the vertex link is a sphere. It follows that none of the 2-quad type solutions is of type B in σ_0 , leaving only the three stated possibilities for the type of σ_0 .

Claim 2: The tetrahedron σ_0 is not of type ACC.

Suppose σ_0 is of type ACC ; we may assume that it is oriented as shown in Figure 11(e). We know that a companion for a square of type A is again of type A . Supposing that the companion square of type A is supported by σ_1 , there are two possible orientations for σ_1 . In either case, a Möbius face of σ_0 has two distinct edges in common with a face of σ_1 . Since the triangulation is face-pair reduced, the faces must be identified. Since face pairings are orientation reversing, this only leaves

one of the two possible orientations of σ_1 . Again applying the fact that the triangulation is face-pair reduced, one finds another pair of faces that is identified. Examining the resulting vertex link shows that it cannot be a sphere; refer to Figure 19.

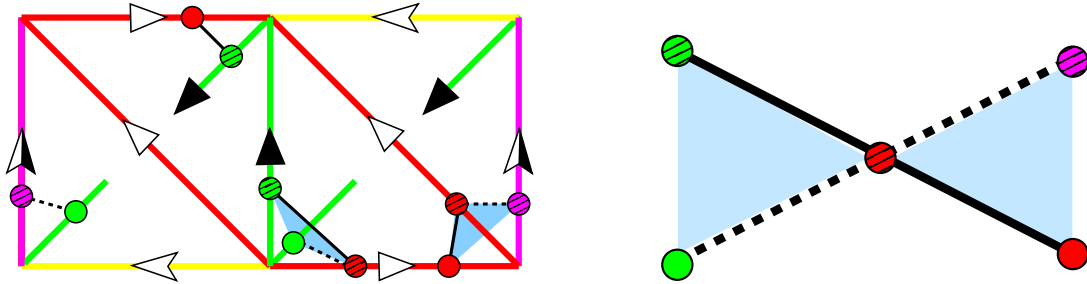


Figure 19: Shown is the only possibility for the companion of a tetrahedron of type ACC. The indicated normal arcs form two simple closed curves that meet transversely in one point on the vertex link. This can be seen by viewing the shaded triangles from the vertex; as the two remaining normal arcs do not intersect in the vertex link, the latter cannot be a sphere.

Claim 3: The tetrahedron σ_0 is not of type AAC.

Suppose σ_0 is of type AAC; we may assume that it is oriented as shown in Figure 11(b). We know that a companion for a square of type C is of type C or D.

First suppose that the companion square is also of type C and supported by σ_1 . Also suppose that one of the squares of type A has companion square supported by σ_1 . For either of the two choices of the square of type A, this forces σ_1 to be of type AAC also. Moreover, σ_0 and σ_1 then have Möbius faces with the same core and boundary edges, and so these faces must be identified. This forces the orientation of σ_1 . Examining the signs of quadrilaterals, it follows that we may also suppose that the companion of the other square of type A is supported by σ_1 . The result of this discussion is shown in Figure 20. Examining the vertex linking triangles contradicts the fact that the vertex link is a sphere.

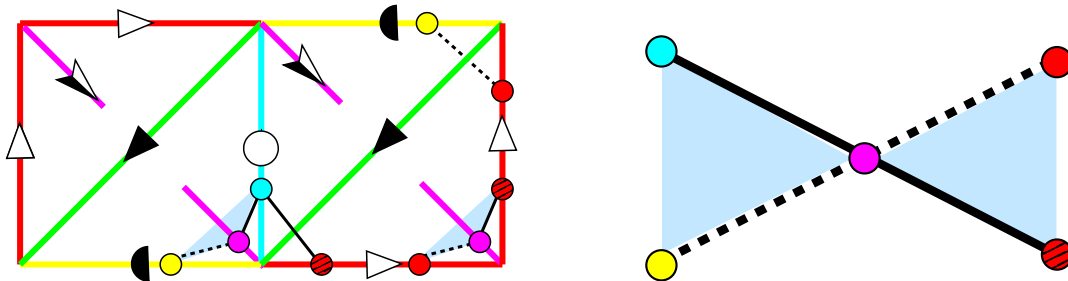


Figure 20: Shown is the first possibility in case AAC. The indicated normal arcs form two simple closed curves that meet transversely in one point on the vertex link. The shaded triangles are viewed from the vertex; as the two additional normal arcs do not intersect in the vertex link, the latter cannot be a sphere.

Hence if the square of type C has a companion square of type C, then one of the squares of type A has companion square supported by $\sigma_2 \neq \sigma_1$, and the other by $\sigma_3 \neq \sigma_1$. If $\sigma_2 \neq \sigma_3$, then the three companion squares are identified as shown in Figure 9(b), contradicting our hypothesis. Hence

$\sigma_2 = \sigma_3$. The union of the two squares of type A in σ_0 is incident with exactly five pairwise distinct edges in M and hence cannot have the same companion square in σ_2 . Moreover, by *Claim 1*, σ_2 must be of type AAC . Examining the signs of quadrilaterals, it follows that we may replace the solution x_1 by a solution x'_1 which is supported by squares of type C in σ_0 and σ_2 . But this is again the same situation of the previous paragraph, where all three solutions are supported in two tetrahedra.

It follows that the companion square for C is of type D . This in particular implies that $\sigma_2 \neq \sigma_1$ and $\sigma_3 \neq \sigma_1$, since a tetrahedron containing a square of type D has no square of type A . If $\sigma_2 = \sigma_3$, then as above, there is a solution x'_1 which is supported by squares of type C in σ_0 and σ_2 , hence putting us into the above case. Whence $\sigma_2 \neq \sigma_3$. The fact that the triangulation is face-pair reduced yields identifications between pairs of faces of the tetrahedra in such a way that not all of them can be orientation reversing; contradicting our hypothesis that the triangulation is oriented (see Figure 21). So AAC cannot happen.

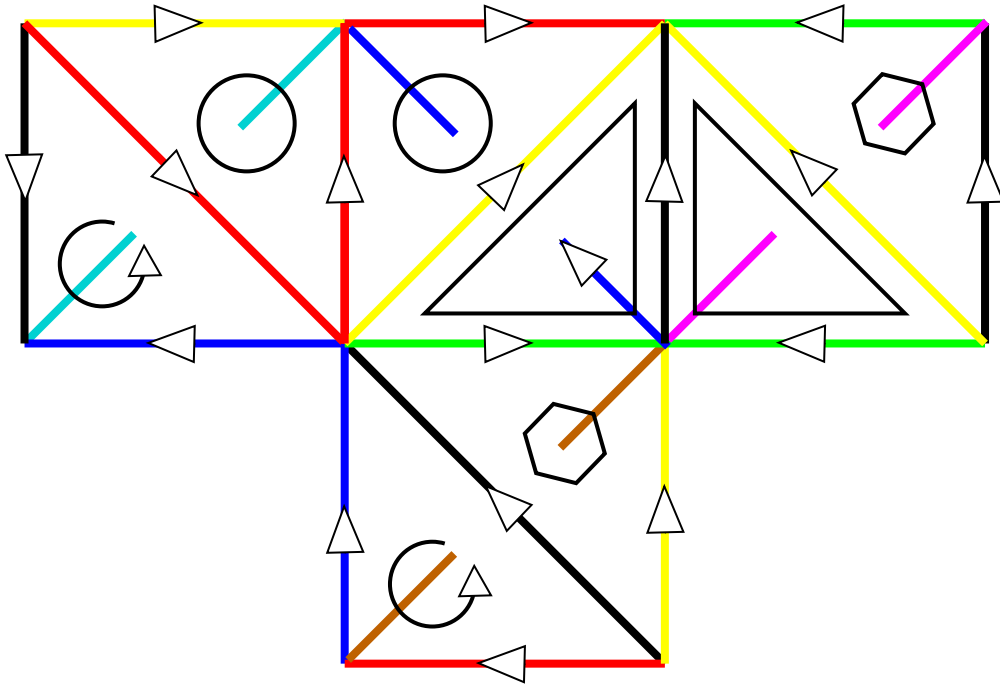


Figure 21: Shown are σ_0 in the middle, with σ_1 on the right, σ_3 on the left and σ_2 below. The edge identifications force identifications of faces, and hence orientations of the tetrahedra. For instance, given an orientation of σ_0 , starting with the pair of faces labelled with a circle determines the orientation of σ_3 . The pair of faces labelled with a triangle determined the orientation of σ_1 . Then the faces labelled with hexagons determine the orientation of σ_2 . But then the faces labelled with a circular arc cannot be identified by an orientation reversing face pairing.

Claim 4: The tetrahedron σ_0 is not of type AAA .

First notice that again we may assume that either $\sigma_1 = \sigma_2 = \sigma_3$, or else the three tetrahedra may be assumed to be pairwise distinct. In case they are pairwise distinct, the three companion squares give the capped projective plane shown in Figure 9(c).

Hence $\sigma_1 = \sigma_2 = \sigma_3$. There are two cases to distinguish depending on the orientation. The first

case cannot occur as the vertex link is a sphere; see Figure 22. In the other case, two squares can be surgered to give a non-separating torus; see Figure 23.

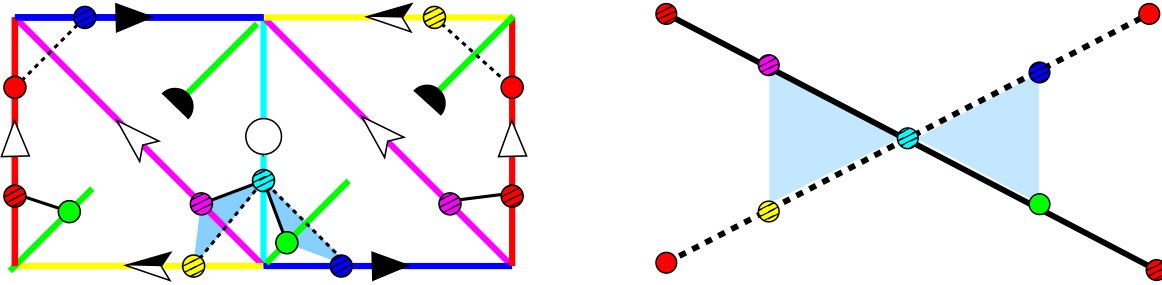


Figure 22: Shown is one possible orientation for the companion of a tetrahedron of type AAA. The indicated normal arcs form two simple closed curves that meet transversely in one point on the vertex link. The transverse intersection can be seen by viewing the shaded triangles from the vertex link.




Figure 23: Shown is the other possible orientation for the companion of a tetrahedron of type AAA. The capped 2-square torus can be viewed as made up of two octagons and two discs on a vertex link. The shown black arc meets the torus transversely and its endpoints can be joined on the vertex link by a path disjoint from the torus; thus certifying that the torus is non-separating.

The four contradictory claims complete the proof of Theorem 23. ■

References

- [1] Benjamin A. Burton: *Face pairing graphs and 3-manifold enumeration*, J. Knot Theory Ramifications 13 (2004), no. 8, 1057–1101.
- [2] Benjamin A. Burton: *Enumeration of non-orientable 3-manifolds using face-pairing graphs and union-find*, Discrete Comput. Geom. 38 (2007), no. 3, 527–571.
- [3] Benjamin A. Burton, Ryan Budney, William Pettersson, et al.: *Regina: Software for 3-manifold topology and normal surface theory*, <http://regina.sourceforge.net/>, 1999–2012.
- [4] William Jaco and J. Hyam Rubinstein: *0-efficient triangulations of 3-manifolds*, Journal of Differential Geometry 65 (2003), no. 1, 61–168.
- [5] William Jaco, J. Hyam Rubinstein and Stephan Tillmann: *Minimal triangulations for an infinite family of lens spaces*, Journal of Topology 2 (2009), no. 1, 157–180.
- [6] William Jaco, J. Hyam Rubinstein and Stephan Tillmann: *\mathbb{Z}_2 -Thurston Norm and Complexity of 3-Manifolds*, Mathematische Annalen 356 (2013) no. 1, 1–22.

- [7] Feng Luo: *Triangulated 3-manifolds: from Haken's normal surfaces to Thurston's algebraic equation*. Interactions between hyperbolic geometry, quantum topology and number theory, 183–204, *Contemp. Math.*, 541, Amer. Math. Soc., Providence, RI, 2011.
- [8] Feng Luo: *Volume Optimization, Normal Surfaces and Thurston's Equation on Triangulated 3-Manifolds*, *Journal of Differential Geometry* 93 (2013), no.2, 299–326.
- [9] Sergei V. Matveev: *Complexity of three-dimensional manifolds: problems and results*, *Siberian Adv. Math.* 13 (2003), no. 3, 95–103.
- [10] Sergei V. Matveev: *Algorithmic topology and classification of 3-manifolds*. Second edition. Algorithms and Computation in Mathematics, 9, Berlin, 2007.
- [11] William P. Thurston: *Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds*. *Ann. of Math. (2)* 124 (1986), no. 2, 203–246.
- [12] Stephan Tillmann, *Normal surfaces in topologically finite 3-manifolds*, *Enseign. Math. (2)* 54 (2008), 329–380.
- [13] Jeffrey L. Tollefson, *Normal surface Q -theory*, *Pacific J. Math.* 183 (1998), no. 2, 359–374.

Feng Luo,
 Department of Mathematics,
 Rutgers University,
 New Brunswick, NJ 08854, USA
 (fluo@math.rutgers.edu)

–

Stephan Tillmann,
 School of Mathematics and Statistics F07,
 The University of Sydney,
 NSW 2006 Australia
 (tillmann@maths.usyd.edu.au)